



# Instrumented ATLAS Athena Jobs on the Grid

**DAN GUNTER, WIM LAVRIJSEN, BRIAN TIERNEY, CRAIG E. TULL**  
**Lawrence Berkeley National Laboratory**  
**National Energy Research Scientific Computing Center**

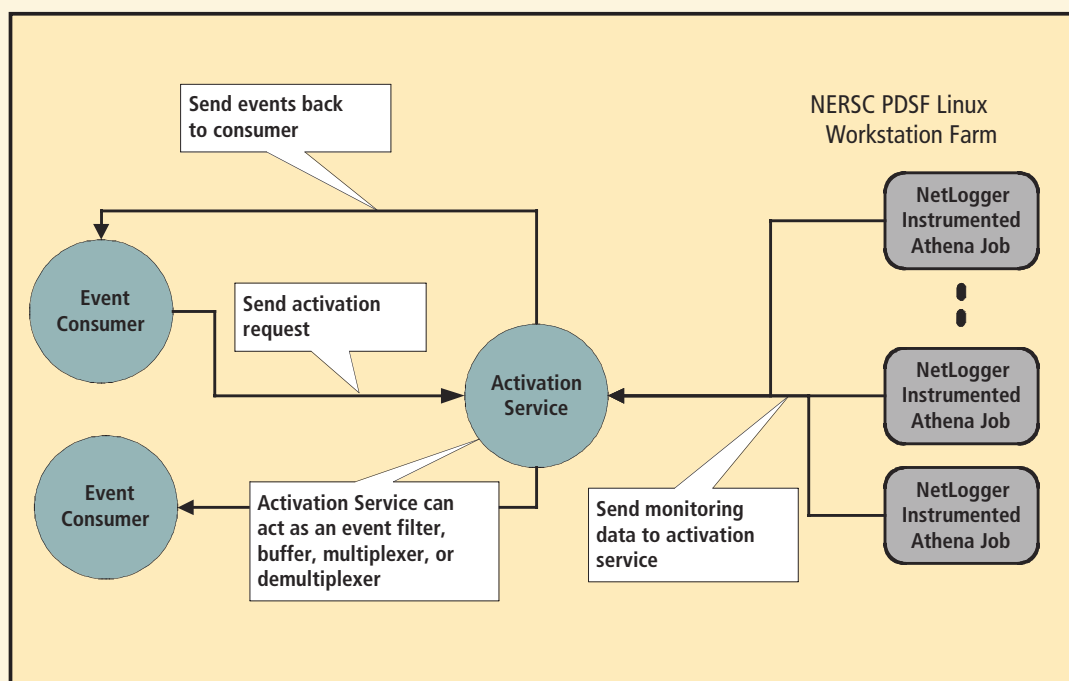
A major goal of the Grid is high-performance distributed computing; but determining whether a Grid application is performing as efficiently as possible requires a considerable amount of instrumentation and monitoring data. To be effective, this monitoring data must be "end-to-end," meaning that components between the Grid application endpoints must also be monitored. This includes all software components (e.g., application software, middleware, operating systems) and all hardware components (e.g., hosts, disks, networks, routers).

Instrumented applications can generate a large amount of monitoring data, so typically the instrumentation is switched off by default. For

jobs running on a Grid, there needs to be a general mechanism to remotely activate the instrumentation in running jobs. The NetLogger Toolkit Activation Service provides this mechanism.

To demonstrate this, we have instrumented the ATLAS Athena Framework with NetLogger to generate monitoring events. We then use a Grid Monitoring Architecture (GMA) based activation service to control NetLogger's trigger mechanism. The NetLogger trigger mechanism allows one to easily start, stop, or change the logging level of a running program by modifying a trigger file. This is shown in the figure below.

A simple GUI allows users to create and modify GMA subscriptions to the activation service. The Activation Service provides a constantly refreshed list of available instrumented applications from which the user can choose. These applications can be on a single cluster, or widely distributed across the Grid.



Instrumentation data from Athena is then sent to the user's desktop, where it can be analyzed using the NetLogger visualization tool, NLV, as shown here.

The Athena object-oriented framework is designed to provide a common infrastructure and environment for simulation, filtering, reconstruction and analysis applications for the current generation of high-energy physics experiments. These experiments are expected to run for many years. Therefore, changes in software requirements and in the technologies used to build software have to be taken into account by developing flexible and adaptable software that can withstand these changes and which can be easily maintained over the long timescales involved. Athena integration with Grid middleware includes the NetLogger instrumentation added by Craig Tull at LBNL. For more information on Athena, see: <http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/00/architecture/General/index.html>

The Activation Service is built using pyGMA, a Python library that makes it easy to develop producers and consumers of monitoring data using the subscription semantics defined by the GMA. GMA is from a Global Grid Forum working group.

NetLogger, pyGMA, and the activation service, are all components being developed as part of the Distributed Monitoring Framework Project at LBNL. For more information see: <http://www.didc.lbl.gov/DMF/>

This work was supported in part by the Director, Office of Science, Office of Basic Energy Research, Division of High Energy Physics of the U.S. Department of Energy under Contract DE-AC03-76SF00098.

